



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/673,755	09/29/2003	Tushar Deepak Chandra	ARC920030059US1	1473
55508 7590 12/07/2007 JOSEPH P. CURTIN, L.L.C. 1469 N.W. MORGAN LANE PORTLAND, OR 97229-5291			EXAMINER PHAM, MICHAEL	
			ART UNIT 2167	PAPER NUMBER
			MAIL DATE 12/07/2007	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/673,755

Applicant(s)

CHANDRA ET AL.

Examiner

Michael D. Pham

Art Unit

2167

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 21 September 2007.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-7,9-18 and 20-22 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-7,9-18 and 20-22 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- ☒ Notice of References Cited (PTO-892)
- ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- ☐ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____
- ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- ☐ Notice of Informal Patent Application
- ☒ Other: _____

Detailed Action

Status of claims

1. Claims 1-7, 9-18, 20-22 are pending.
2. Claims 1-7, 9-18, 20-22 have been examined.

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. **Claims 1 - 7 and 12-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Watt (US Patent Application Publication 2003/0126202; hereafter Watt) further in view of Sheets et al. (US 6,816,905; hereafter Sheets).**

Claim 1:

Watt discloses the following claimed limitations:

“identifying at least one master storage image that is stored in the storage of the storage area network and that will be associated with a system user when a server is allocated to the system user;” [0046, the administrator can pick and choose from the installed software base to create a master server image. And that the repository manager can also install and manage instances of a SAN as well as on a server’s local attached storage.]

“generating a plurality of replicas of each identified master storage image prior to at least one server being allocated to the system user; and” [0046, once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.]

“allocating a selected replica of the plurality of replicas of the master storage image to each server allocated to the system user.” [0046, once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.]

Watt’s does not explicitly disclose “pre-configuring at least one identified master storage image with data and state information that is associated with a system user”.

Sheets disclosed col. 7 lines 1-5, collects status information from the other servers and dynamically reallocates at least one server from the first administrative group to the second administrative group in response to at least the status information. Further disclosing, abstract, the automatic and dynamic allocation of servers is accomplished by setting initialization pointers for the reallocated servers to access software and data unique to the customer account for the second administrative group, and then reinitializing the reallocated servers such that the reallocated servers join the second administrative group when restarted. col. 15 lines 8-17, one of the significant advantages of the present invention is the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative

group. This provides a natural and very efficient security mechanism for precluding intentional or unintentional access to data between different customer accounts. Accordingly, pre-configuring (col. 15 lines 8-17, re-configuring) at least one identified master storage image (abstract, administrative group) with data (abstract, data) and state information (col. 7 lines 1-5, status) that is associated with a system user (abstract, customer) is suggested.

Both Watts and Sheets are within the same field of endeavor as the application, namely dynamic allocation of servers. Watts suggests configuring images associated with a user. Sheets suggests pre-configuring at least one identified server image with data and state information associated with a user. It would have been obvious to apply Sheets disclosure to Watts to improve security by keeping the data separate from each user, and further to provide a method of indicating states of servers in order to improve allocation techniques. Therefore, improving performance of server allocation as a whole, and further improving security as the unique data can be kept separate by precluding intentional or unintentional access to data between different customer accounts.

Claim 2:

Watt does not explicitly disclose “de-allocating an allocated replica from the system user each time a server is de-allocated from the system user; and assigning the de-allocated replica to a pool of de-allocated replicas.”

However, Sheets discloses “de-allocating an allocated replica from the system user each time a server is de-allocated from the system user; and assigning the de-allocated replica to a

Art Unit: 2167

pool of de-allocated replicas". (See column 18, lines 32-45, discloses, another way of looking at how the present invention can dynamically provide hosted service across disparate accounts is to view a portion of the servers as being assigned to a pool of a plurality of virtual servers that may be selectively configured to access software and data for a particular administrative group. Further disclosing it automatically allocates one of the servers from the pool of virtual servers to that administrative group. Conversely, if the dispatch module determines that an agency group can relinquish one of its servers, that relinquished server would be added to the pool of virtual servers that are available for re allocation to a different administrative group.)

It would have been obvious to one with ordinary skill in the art to combine the method as disclosed in Watt with the de-allocating method as disclosed in Sheets et al. because the two both disclose methods that operate very similarly, but the de-allocation was simply not addressed in Watt, but is a natural extension of Watt. By moving the de-allocated replica into the pool, the server is now marked available for future use when it is needed. It is for this reason that one of ordinary skill in the art would have been motivated to include de-allocating an allocated replica from the system user each time a server is de-allocated from the system user; and assigning the de-allocated replica to a pool of de-allocated replicas.

Claim 3:

Watt does not explicitly disclose "the pool of de-allocated replicas is configured to automatically scrub all replicas in the pool of de-allocated replicas asynchronously from de-allocation the step of de-allocation."

However, Sheets discloses “the pool of de-allocated replicas is configured to automatically scrub all replicas in the pool of de-allocated replicas asynchronously from de-allocation the step of de-allocation” (See column 15, lines 8-14, discloses one of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.)

It would have been obvious to one with ordinary skill in the art to combine the method of Watt with the automatic scrub function of Sheets because of the advantage automatically clearing the unique data provides. This provides a layer of security, allowing different users to use the same equipment dynamically without having to worry about improper sharing of secret data. It is for this reason that one of ordinary skill in the art would have been motivated to have the pool of de-allocated replicas configured to automatically scrub all replicas in the pool of de-allocated replicas asynchronously from de-allocation the step of de-allocation.

Claim 4:

Watt does not explicitly disclose “the pool of de-allocated replicas is scrubbed when a number of de-allocated replicas assigned to the pool of de-allocated replicas equals a predetermined number.”

However, Sheets discloses “the pool of de-allocated replicas is scrubbed when a number of de-allocated replicas assigned to the pool of de-allocated replicas equals a predetermined number.” [See column 15, lines 8-14, discloses that one of the significant advantages of the

present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.]

It would have been obvious to one with ordinary skill in the art to combine the method of Watt with the automatic scrub function of Sheets because of the advantage automatically clearing the unique data provides. This provides a layer of security, allowing different users to use the same equipment dynamically without having to worry about improper sharing of secret data. While not specifically mentioned here, the predetermined number can be considered to be one and the replica is scrubbed by reconfiguring it for use with another user's data. It is for this reason that one of ordinary skill in the art would have been motivated to have the pool of de-allocated replicas is scrubbed when a number of de-allocated replicas assigned to the pool of de-allocated replicas equals a predetermined number.

Claim 5:

Watt does not explicitly disclose "the pool of de-allocated replicas is automatically scrubbed by reformatting."

However, Sheets discloses "the pool of de-allocated replicas is automatically scrubbed by reformatting." (See column 15, lines 8-14, discloses one of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer

account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.)

It would have been obvious to one with ordinary skill in the art to combine the method of Watt with the automatic scrub by reformatting function of Sheets because of the advantage automatically clearing the unique data provides. This provides a layer of security, allowing different users to use the same equipment dynamically without having to worry about improper sharing of secret data. It is for this reason that one of ordinary skill in the art would have been motivated to have the pool of de-allocated replicas is automatically scrubbed by reformatting.

Claim 6:

Watt additionally discloses “each replica is a logical volume.” (0067 discloses that the SAN routing and volume assignment can be changed by DSAP system 102 thereby affecting the SAN's mapping of the server's SAN connection to a SAN volume. The way the replica is referred to here makes it clear that the replicas are logical volumes.)

Claim 7:

Watt does not explicitly disclose “the system user is one of a customer and an application”. However, Sheets discloses that the system user is one of a customer and an application. (See column 15, lines 8-14, discloses one of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought

into service as part of the second administrative group.) It would have been obvious to one with ordinary skill in the art to combine the method of Watt with the automatic scrub by reformatting function of Sheets because of the advantage automatically clearing the unique data provides. This provides a layer of security, allowing different users to use the same equipment dynamically without having to worry about improper sharing of secret data. It is also clear that the user is referring to a customer in Sheets. It is for this reason that one of ordinary skill in the art would have been motivated to have the system user is one of a customer and an application.

Claim 12:

Watt discloses the following claimed limitations:

“a plurality of servers coupled to a storage” [0046, repository manager is responsible for securely and efficiently provisioning and managing server images on storage devices within data centers. And figure 2 elements 212, 218, and 210.]; and

“a storage provisioning device coupled to the servers and allocating at least one server and a portion of the storage to a system user, the storage provisioning device identifying at least one master storage image that is stored in the storage and that will be associated with a system user when a server is allocated to the system user” [0046, the administrator can pick and choose from the installed software base to create a master server image. 0058, in an alternate embodiment, repository manager can also install and manage instances of a SAN as well as on a server’s local attached storage.]

“the storage provisioning device further generating a plurality of replicas of each identified master storage image prior to at least one server being allocated to the system user”

[0046, once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool]; and

“allocating a selected replica of the plurality of replicas of the master storage image to each server allocated to the system user.”[0046, once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool]

Watt does not explicitly disclose “at least one master storage image being pre-configured with data and state information that is associated with a system user”.

Sheets disclosed col. 7 lines 1-5, collects status information from the other servers and dynamically reallocates at least one server from the first administrative group to the second administrative group in response to at least the status information. Further disclosing, abstract, the automatic and dynamic allocation of servers is accomplished by setting initialization pointers for the reallocated servers to access software and data unique to the customer account for the second administrative group, and then reinitializing the reallocated servers such that the reallocated servers join the second administrative group when restarted. col. 15 lines 8-17, one of the significant advantages of the present invention is the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group. This provides a natural and very efficient security mechanism for precluding intentional or unintentional access to data between different customer accounts. Accordingly, pre-

configuring (col. 15 lines 8-17, re-configuring) at least one identified master storage image (abstract, administrative group) with data (abstract, data) and state information (col. 7 lines 1-5, status) that is associated with a system user (abstract, customer) is suggested.

Both Watts and Sheets are within the same field of endeavor as the application, namely dynamic allocation of servers. Watts suggests configuring images associated with a user. Sheets suggests pre-configuring at least one identified server image with data and state information associated with a user. It would have been obvious to apply Sheets disclosure to Watts to improve security by keeping the data separate from each user, and further to provide a method of indicating states of servers in order to improve allocation techniques. Therefore, improving performance of server allocation as a whole, and further improving security as the unique data can be kept separate by precluding intentional or unintentional access to data between different customer accounts.

Claim 13-17:

Regarding claim 13-18, Watt in view of Sheets disclose the storage area network as cited above for claims 2-7 respectively.

5. Claims 9 and 20 are rejected under 35 U.S.C. 103(a) as being unpatentable over Watt (US Patent Application Publication 2003/0126202; hereafter Watt) further in view of Haun et al. (US 6,751,658; hereafter Haun).

Claim 9:

Watts discloses the following claimed limitations:

“identifying at least one master storage image that is stored in the storage of the storage area network and that will be associated with a system user, each master storage image including both a read-only data portion and a writeable data portion;” [0046, the administrator and pick and choose from the installed software base to create a master server image. and 0058, in an alternate embodiment, repository manager can also install and manage instances of a SAN as well as on a server’s local attached storage.]

“generating a read-only copy of the read-only data portion of each master storage image;” [0095, an independent instance contains an actual physical copy of all files in the master image, with the configuration files updated to provide a unique personally. The independent instance is stored on centralized storage and can be run by any available server.]

“sharing the read-only data copy of the read-only data portion of each master storage image across the plurality of servers;” [0097, the remainder of the image is shared with other dependent instances by referencing the read-only snapshot containing the original files]

“allocating the read-only copy of the read-only data portion of a selected master storage image to each server allocated to the system user; and” [0046, once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.”]

However, Watts does not explicitly disclose, “allocating a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user.”

On the other hand Haun, Abstract, discloses a network computer (NC) client boots from a boot image provided by an NC server. The boot image includes information identifying the location of one or more system volumes on the NC server that contain operating system software. In response to an attempt to modify the contents of the one or more system volumes, the NC client causes information identifying the modification to be recorded on the NC server separate from the one or more system volumes in a storage area associated with the NC client. figure 4, element 430, the NC server stores information associated with the write request in the shadow volume associated with the NC client. Accordingly, allocating a separate writable volume of the writable portion of the selected master storage image to each server (modification to be recorded on the NC server) allocated to the system user (client) is suggested.

Both Watt and Haun are within applicant's same field of endeavor as they are directed towards provisioning computer systems. It would have been obvious to a person of an ordinary skill in the art to have applied Haun's disclosure above, to the system of Watt for the purpose of maintaining individual user storage areas. In doing so, it provides the benefit of easily upgrading individual systems. It is for the above reasons that one of ordinary skill in the art at the time the invention was made would have been motivated to allocate a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user.

Claim 20:

Watts discloses the following claimed limitations:

“a plurality of servers coupled to a storage; and” [0046, repository manager is responsible for securely and efficiently provisioning and managing server images on storage devices within data centers. And figure 2 elements 212, 218, and 210.]

“a storage provisioning device coupled to the servers and allocating at least one server and a portion of the storage to a system user, the storage provisioning device identifying at least one master storage image that is stored in the storage of the storage area network and that will be associated with a system user, each master storage image including both a read-only data portion and a writeable data portion,” [0046, the administrator and pick and choose from the installed software base to create a master server image. and 0058, in an alternate embodiment, repository manager can also install and manage instances of a SAN as well as on a server’s local attached storage.]

“the storage provisioning device further generating a read-only copy of the read-only portion of each master storage image” [0095, an independent instance contains an actual physical copy of all files in the master image, with the configuration files updated to provide a unique personally. The independent instance is stored on centralized storage and can be run by any available server.]

“and sharing the read-only copy of the read-only portion of each master storage image across the plurality of servers,” [0097, the remainder of the image is shared with other dependent instances by referencing the read-only snapshot containing the original files]

“allocating the read-only copy of the read-only portion of a selected master storage image to each server allocated to the system user,” [0046, once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.”]

However, Watts does not explicitly disclose, “allocating a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user.”

On the other hand Haun, Abstract, discloses a network computer (NC) client boots from a boot image provided by an NC server. The boot image includes information identifying the location of one or more system volumes on the NC server that contain operating system software. In response to an attempt to modify the contents of the one or more system volumes, the NC client causes information identifying the modification to be recorded on the NC server separate from the one or more system volumes in a storage area associated with the NC client. figure 4, element 430, the NC server stores information associated with the write request in the shadow volume associated with the NC client. Accordingly, allocating a separate writable volume of the writable portion of the selected master storage image to each server (modification to be recorded on the NC server) allocated to the system user (client) is suggested.

Both Watt and Haun are within applicant’s same field of endeavor as they are directed towards provisioning computer systems. It would have been obvious to a person of an ordinary

skill in the art to have applied Haun's disclosure above, to the system of Watt for the purpose of maintaining individual user storage areas. In doing so, it provides the benefit of easily upgrading individual systems. It is for the above reasons that one of ordinary skill in the art at the time the invention was made would have been motivated to allocate a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user.

6. Claims 10 – 11 and 21 - 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Watt (US Patent Application Publication 2003/0126202; hereafter Watt) further in view of Haun et al. (US 6,751,658; hereafter Haun) and Sheets et al. (US 6,816,905; hereafter Sheets).

Claim 10:

Watt and Haun disclose a method substantially as claimed.

Watt and Haun do not explicitly disclose “de-allocating the read-only copy of the read-only data portion of the selected master image from the server to which the read-only copy was allocated when the server is de-allocated from the system user; and de-allocating the writable data volume of the writable data portion of the selected master storage image that was allocated to the de-allocated server.”

However, Sheets discloses de-allocating the read-only copy of the read-only data portion of the selected master image from the server to which the read-only copy was allocated when the server is de-allocated from the system user; and de-allocating the writable data volume of the

writable data portion of the selected master storage image that was allocated to the de-allocated server. (See column 18, lines 32-45 discloses another way of looking at how the present invention can dynamically provide hosted service across disparate accounts is to view a portion of the servers as being assigned to a pool of a plurality of virtual servers that may be selectively configured to access software and data for a particular administrative group...it automatically allocates one of the servers from the pool of virtual servers to that administrative group. Conversely, if the dispatch module determines that an agency group can relinquish one of its servers, that relinquished server would be added to the pool of virtual servers that are available for re allocation to a different administrative group.)

It would have been obvious to one with ordinary skill in the art to combine the method as disclosed in Watt and Haun with the de-allocating method as disclosed in Sheets because they disclose methods that operate very similarly, but the de-allocation was simply not addressed in Watt and Haun, but is a natural extension of Watt and Haun. By moving the de-allocated copy into the pool, the server is now marked available for future use when it is needed. It is for this reason that one of ordinary skill in the art would have been motivated to include de-allocating the read-only copy of the read-only data portion of the selected master image from the server to which the read-only copy was allocated when the server is de-allocated from the system user; and de-allocating the writable data volume of the writable data portion of the selected master storage image that was allocated to the de-allocated server.

Claim 11:

Watt and Haun discloses a method substantially as claimed.

Watt and Haun do not explicitly disclose “de-allocating the writable data volume includes the steps of: assigning the de-allocated writable data volume to a pool of de-allocated writable data volumes; and scrubbing any writable data volumes assigned to the pool of de-allocated writable data volumes asynchronously from the step of de-allocating the writable data volume.”

However, Sheets suggests “de-allocating the writable data volume includes the steps of: assigning the de-allocated writable data volume to a pool of de-allocated writable data volumes; and scrubbing any writable data volumes assigned to the pool of de-allocated writable data volumes asynchronously from the step of de-allocating the writable data volume” [column 18, lines 41-45 discloses conversely, if the dispatch module determines that an agency group can relinquish one of its servers, that relinquished server would be added to the pool of virtual servers that are available for re allocation to a different administrative group. Further disclosing column 15, lines 8-14 one of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.]

It would have been obvious to one with ordinary skill in the art to combine the teachings of Watt and Haun with the disclosure in Sheets by adding the scrubbing method to enhance the security of sharing the data volumes between different users. It is for this reason that one of ordinary skill in the art would have been motivated to have the step of de-allocating the writable data volume include the steps of: assigning the de-allocated writable data volume to a pool of de-allocated writable data volumes; and scrubbing any writable data volumes assigned to the pool of

Art Unit: 2167

de-allocated writable data volumes asynchronously from the step of de-allocating the writable data volume.

Claim 21-22:

Regarding claim 21 and 22, Watt and Haun in view of Sheets disclose the storage area network as cited above for claims 10 and 11 respectively.

Response to Arguments

7. Applicant's arguments filed 9/21/07 have been fully considered but they are not persuasive. Applicant's have asserted the following (lettered):

A. There is a prima facie case of obviousness is not met because no motivation is given for the combination of Watt and Sheets to be combined and is using impermissible hindsight reasoning (remarks pages 7-11).

In response to Applicant's assertions that a prima facie case of obviousness is not met, the examiner respectfully disagrees that there is no motivation to combine Watts and Sheets.

In regards to Applicant's assertions page 8, that the motivation must come from Sheets. The examiner respectfully disagrees. It is clear, from Applicant's statement on page 7 that this should be understood. Applicant's clearly agree that "...there must be some suggestion or motivation, either in the references themselves or the knowledge generally available to one of ordinary skill in the art to modify the references or to combine references teachings."

Furthermore, the motivation is provided in Sheets. Sheets states in col. 15 lines 15-17, that this would provide a natural and very efficient security mechanism for precluding intentional or unintentional access to data between different customer accounts. Watts paragraph 0020 states that many network security issues are eliminated by automatically configuring the network infrastructure when a server is provisioned to restrict access to just those resources within the data center that the server needs to perform its function. It would have been obvious to apply Sheets disclosure to Watts to improve security by keeping the data separate from each user, and further to provide a method of indicating states of servers in order to improve allocation techniques. Therefore, improving performance of server allocation as a whole, and further improving security as the unique data can be kept separate by precluding intentional or unintentional access to data between different customer accounts.

As to applicant's assertions directed towards page 8, that the examiner still does not indicate that Watts suggests that the master storage image can be pre-configured with data and state information that is associated with a system user. In response, the examiner stated this limitation was not explicitly disclosed by Watts. However indicated that, Watts still disclosed configuring images associated with a user. As it is clearly stated in 0046, the administrator (associated with user) can pick and choose from the installed software base to create (configuring) a master server image (master server image). Once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool. Watt however does not explicitly disclose "pre-configuring at least one identified master storage image with data and state information that is associated with a system user".

On the other hand, Sheets disclosed col. 7 lines 1-5, collects status information from the other servers and dynamically reallocates at least one server from the first administrative group to the second administrative group in response to at least the status information. Further disclosing, abstract, the automatic and dynamic allocation of servers is accomplished by setting initialization pointers for the reallocated servers to access software and data unique to the customer account for the second administrative group, and then reinitializing the reallocated servers such that the reallocated servers join the second administrative group when restarted. col. 15 lines 8-17, one of the significant advantages of the present invention is the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group. This provides a natural and very efficient security mechanism for precluding intentional or unintentional access to data between different customer accounts. Accordingly, pre-configuring (col. 15 lines 8-17, re-configuring) at least one identified master storage image (abstract, administrative group) with data (abstract, data) and state information(col. 7 lines 1-5, status) that is associated with a system user (abstract, customer) is suggested.

It would have been obvious to apply Sheets disclosure to Watts to improve security by keeping the data separate from each user, and further to provide a method of indicating states of servers in order to improve allocation techniques. Therefore, improving performance of server allocation as a whole, and further improving security as the unique data can be kept separate by precluding intentional or unintentional access to data between different customer accounts.

In response to applicant's argument that the examiner's conclusion of obviousness is based upon improper hindsight reasoning, it must be recognized that any judgment on obviousness is in a sense necessarily a reconstruction based upon hindsight reasoning. But so long as it takes into account only knowledge which was within the level of ordinary skill at the time the claimed invention was made, and does not include knowledge gleaned only from the applicant's disclosure, such a reconstruction is proper. See In re McLaughlin, 443 F.2d 1392, 170 USPQ 209 (CCPA 1971). In this case, Watts discloses configuring master images associated with a system user, as noted above. Sheets discloses more explicitly pre-configuring the master images with data and state information associated with the system user, as noted above. Accordingly, it would have been obvious to a person of an ordinary skill to utilize the data and state information in Sheets to the disclosure of Watts for the purpose of determining the status of servers and dynamically reallocate servers according to the customer needs. By using the method of Sheets it further provides the advantage of precluding intentional and unintentional access to customer accounts. Thus, adding to Watts' restricting access for security measures, and further Watts' monitoring of servers. Ergo, Watts and Sheets are directed to similar goals, and further relate to applicant's field of endeavor. Accordingly, the combination of the two references would be well within the level of a person of an ordinary skill. Therefore, applicant's assertions directed towards the combination of Watts and Sheets are unpersuasive.

B. There is no motivation for Watt and Haun to be combined (remarks pages 12-14).

In response to Applicant's assertions that a prima facie case of obviousness is not met, the examiner respectfully disagrees that there is no motivation to combine Watts and Haun.

In regards to Applicant's assertions page 13, that the motivation must come from Sheets. The examiner respectfully disagrees. It is clear, from Applicant's statement on page 12 that this should be understood. Applicant's clearly agree that "...there must be some suggestion or motivation, either in the references themselves or the knowledge generally available to one of ordinary skill in the art to modify the references or to combine references teachings."

In regards to Applicant's assertion directed towards that there is nothing in Haun that indicates the limitation "allocating a separate writable data volume of the writable portion of the selected master storage image to each server allocated to the system user". The examiner respectfully disagrees. Haun, Abstract, discloses a network computer (NC) client boots from a boot image provided by an NC server. The boot image includes information identifying the location of one or more system volumes on the NC server that contain operating system software. In response to an attempt to modify the contents of the one or more system volumes, the NC client causes information identifying the modification to be recorded on the NC server separate from the one or more system volumes in a storage area associated with the NC client. figure 4, element 430, the NC server stores information associated with the write request in the shadow volume associated with the NC client. Accordingly, allocating a separate writable volume of the writable portion of the selected master storage image to each server (modification to be recorded on the NC server) allocated to the system user (client) is suggested.

Both Watt and Haun are within applicant's same field of endeavor as they are directed towards provisioning computer systems. It would have been obvious to a person of an ordinary skill in the art to have applied Haun's disclosure above, to the system of Watt for the purpose of maintaining individual user storage areas. In doing so, it provides the benefit of easily upgrading individual systems. It is for the above reasons that one of ordinary skill in the art at the time the invention was made would have been motivated to allocate a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user.

In regards to applicant's assertions stating that the writable volume referred to in the portions of Haun cited by examiner might actually be a single writable volume. Haun discloses multiple volumes also see col. 4 lines 35-41.

*In response to applicant's argument that the examiner's conclusion of obviousness is based upon improper hindsight reasoning, it must be recognized that any judgment on obviousness is in a sense necessarily a reconstruction based upon hindsight reasoning. But so long as it takes into account only knowledge which was within the level of ordinary skill at the time the claimed invention was made, and does not include knowledge gleaned only from the applicant's disclosure, such a reconstruction is proper. See *In re McLaughlin*, 443 F.2d 1392, 170 USPQ 209 (CCPA 1971). In this case, Watts provides a system that provisions servers. Haun, Abstract, discloses a network computer (NC) client boots from a boot image provided by an NC server. The boot image includes information identifying the location of one or more*

Art Unit: 2167

system volumes on the NC server that contain operating system software. In response to an attempt to modify the contents of the one or more system volumes, the NC client causes information identifying the modification to be recorded on the NC server separate from the one or more system volumes in a storage area associated with the NC client. figure 4, element 430, the NC server stores information associated with the write request in the shadow volume associated with the NC client. Accordingly, allocating a separate writable volume of the writable portion of the selected master storage image to each server (modification to be recorded on the NC server) allocated to the system user (client) is suggested. Both Watt and Haun are within applicant's same field of endeavor as they are directed towards provisioning computer systems. It would have been obvious to a person of an ordinary skill in the art to have applied Haun's disclosure above, to the system of Watt for the purpose of maintaining individual user storage areas. In doing so, it provides the benefit of easily upgrading individual systems. It is for the above reasons that one of ordinary skill in the art at the time the invention was made would have been motivated to allocate a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user. Accordingly, the combination of the two references would be well within the level of a person of an ordinary skill. Therefore, applicant's assertions directed towards the combination of Watts and Sheets are unpersuasive.

C. Claims 10, 11, 21, and 22 are patentable based on the remarks provided for claims 9 and 20 (remarks page 15).

In response, see B.

Conclusion

8. The prior art made of record listed on PTO-892 and not relied, if any, upon is considered pertinent to applicant's disclosure.

Contact Information

9. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael D. Pham whose telephone number is (571)272-3924. The examiner can normally be reached on Monday - Friday 9am - 5:00pm.

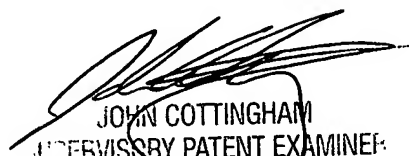
If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Cottingham can be reached on 571-272-7079. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Michael Pham
Art Unit 2167
Examiner *m.p.*

Cam Y. Truong
Art Unit 2162
Primary Examiner
cy

John Cottingham
Art Unit 2167
Supervisor


JOHN COTTINGHAM
INTERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100